


May 1991

NASA-TP-3089 19910016427

Model Reduction by Trimming for a Class of Semi-Markov Reliability Models and the Corresponding Error Bound

Allan L. White
and Daniel L. Palumbo



1991

**Model Reduction by Trimming
for a Class of Semi-Markov
Reliability Models and the
Corresponding Error Bound**

Allan L. White
and Daniel L. Palumbo
Langley Research Center
Hampton, Virginia



National Aeronautics and
Space Administration
Office of Management
Scientific and Technical
Information Division

Abstract

Semi-Markov processes have proved to be an effective and convenient tool for constructing models of systems that achieve reliability by redundancy and reconfiguration. These models are able to depict complex system architectures and to capture the dynamics of fault arrival and system recovery. A disadvantage of this approach is that the models can be extremely large, which poses both a model construction and a computational problem. Techniques are needed to reduce the model size. Because these systems are used in critical applications where failure can be expensive, there must be an analytically derived bound for the error produced by the model reduction technique. This report presents a model reduction technique called trimming that can be applied to a popular class of systems. Automatic model generation programs have been written to help the reliability analyst produce models of complex systems. This method (trimming) is easy to implement and its error bound easy to compute. Hence, the method lends itself to inclusion in an automatic model generator.

Introduction

Reliable digital control systems are being designed using redundancy and reconfiguration. The reliability requirement for these systems can be extremely high. An example is the proposed requirement that the flight control system for a commercial aircraft have less than one chance in a billion of failure during a 10-hour flight. Such requirements are beyond what can be established by natural life testing. An alternative method is to estimate the probability of system failure with a semi-Markov model that captures the elements of system architecture, component failure, and system recovery from failed components. The system architecture can be described by considering the components and how they interact. The component failure rate is obtained from field data. The description of system recovery from failed components is determined from fault-injection experiments in the laboratory. These three features (system architecture, component failure, and system recovery) can be studied separately and then combined to form the reliability model for the system. Semi-Markov processes with their states representing the states of the system and their transitions between states representing fault occurrences and system recoveries have proved to be an effective and convenient reliability estimation tool.

A major obstacle is that a reconfigurable system of moderate size and complexity can generate an enormous semi-Markov model, producing both a model construction problem and a computational problem. The model construction problem has become severe enough that computer programs have been written to automatically generate reliability

models. These automatic model generators have intensified the computational problem, since it is now possible to produce models of extremely complex systems.

Sound and effective procedures are needed for model reduction. Since these models describe systems that need to be highly reliable, an acceptable model reduction method must have an analytically derived error bound. Since the model reduction method presented in this report is easy to implement and its error bound easy to compute, it lends itself to inclusion in an automatic model generator. In fact, it is currently being developed as a feature of the automatic model generator called ASSIST (ref. 1).

We call the procedure model reduction by trimming, or just trimming (ref. 2). In the next section, we illustrate trimming by means of a concrete example. Then in subsequent sections trimming is precisely defined, the theorem for the error bound on trimming is precisely stated, and the trimming bound is derived. Finally we show that not all models can be trimmed and still yield an accurate estimate of reliability. It is essential to determine the error produced by trimming.

Illustrative Example

This section uses a simple example to illustrate the basic ideas of model reduction by trimming. This example is not completely realistic in engineering terms, but it covers the ideas in a concrete manner. Suppose a system consists of four central processor units, four memories, and four buses. In the initial configuration, three components of each type are active while the fourth is a cold spare (with zero failure

rate). If a component becomes faulty, it is replaced by a spare. If the number of good processors, good memories, or good buses falls below three, then the entire system goes into a simplex configuration consisting of one processor, one memory, and one bus. In this simplex configuration, the failure of any of the three components causes system failure. The initial configuration, showing only the active components, is displayed in figure 1. In this initial configuration, and in the subsequent configurations as a triad, each processor (CPU) is assigned one bus to use for sending data to all three memories. Each memory (MEM) receives data from all three processors. Similarly, each memory is assigned one bus to use for sending data to all three processors. Each processor receives data from all three memories.

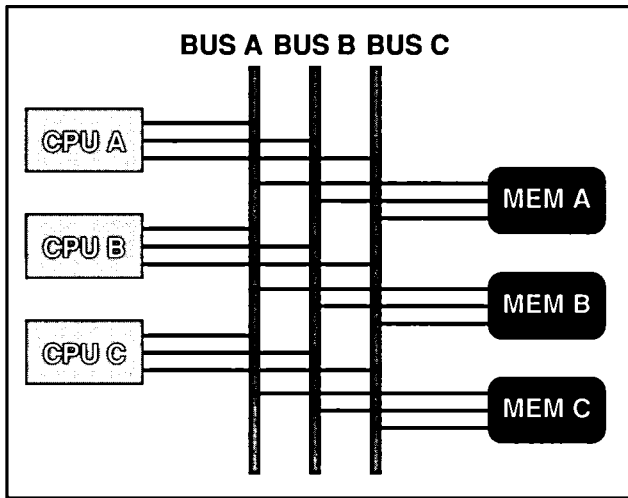


Figure 1. Initial configuration of processors, memories, and buses.

The system begins an operation cycle with the active processors requesting data from the memories. Each memory (on its assigned bus) sends data to all three processors. Each processor votes on the received data (i.e., the data from the three memories are compared to detect a fault should the data disagree) and performs its calculations. After computing, each processor (on its assigned bus) sends data to all three memories. An operation cycle ends when each memory votes and stores the data.

Some critical coupling exists between the processors and buses in the sense that the system can have a coincident-fault failure when one fault is in a processor and the other fault is in a bus. For example, suppose the processors are sending data to the memories with processor i using bus i for $i = A, B, C$. If processor A and bus B are faulty, then the memory voters can be overwhelmed by incorrect data. However, if processor A and bus A are faulty, then the

memories will vote correctly. There is similar critical coupling between the memories and the buses. There is no critical coupling between the processors and the memories.

Part of the reliability model for this system is shown in figure 2. The failure rates for processors, memories, and buses are λ_P , λ_M , and λ_B , respectively. For convenience, the system recoveries are assumed to be constant-rate transitions with F_P , F_M , and F_B being the system recovery rates for processors, memories, and buses. The states are denoted by A for a fault-free state, R for a single-fault recovery-mode state, V for a multiple-fault recovery-mode state, and X and Y for system failure states.

To illustrate model reduction by trimming, consider state R_P of figure 2 where one of the active processors has become faulty. The recovery transition F_P removes this faulty processor and replaces it with the spare. The transition $2\lambda_P + 2\lambda_B$ represents the failure of another processor or of a bus that is critically coupled to the failed processor. The state X represents system failure because of coincident faults. The transitions $2\lambda_M$, λ_M , and λ_B represent fault arrival in components that are not critically coupled to the faulty processor. It seems reasonable to think that these last three transitions and their subsequent states can be ignored with negligible loss of accuracy, because even after these transitions there must be another component failure before there is system failure. We call such states as R_P recovery-mode states. A recovery-mode state is a state with a recovery transition out of it. Model reduction by trimming eliminates all component failure transitions from recovery-mode states that do not cause immediate system failure.

Two models, complete and trimmed, were constructed for this system using the ASSIST reliability model generator (ref. 3). The complete model contains 227 states and took 7878 cpu (central processing unit) seconds to compute on a Digital Equipment Corp. VAX 11/750 computer. The trimmed model contains 83 states and took 258 cpu seconds to compute on the same computer. Different methods of constructing a reliability model can produce equivalent models with different numbers of states, but the relative difference between the complete and the trimmed model is thought to remain the same.

For the parameter values,

$$\lambda_P = 10^{-4}/\text{hour} \quad F_P = 10^4/\text{hour}$$

$$\lambda_M = 5 \times 10^{-4}/\text{hour} \quad F_M = 10^3/\text{hour}$$

$$\lambda_B = 10^{-5}/\text{hour} \quad F_B = 10^3/\text{hour}$$

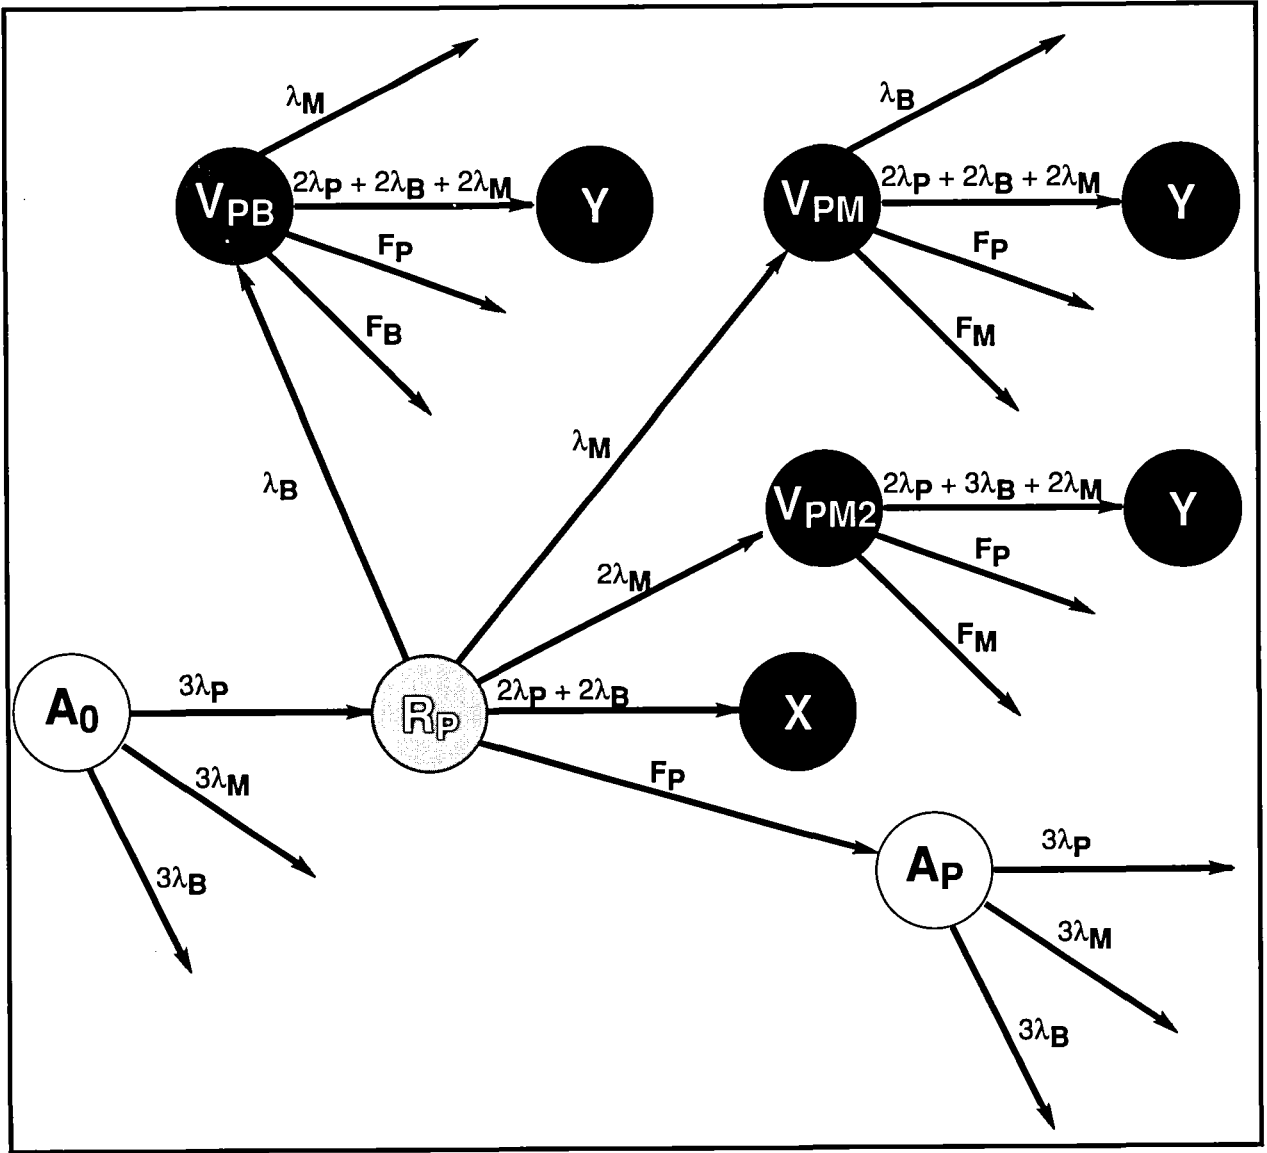


Figure 2. Reliability model for the illustrative example in figure 1.

and an operating time of $T = 1$ hour, the complete model returns for the probability of system failure

$$P(\text{Failure of complete model}) = 1.80803726 \times 10^{-9}$$

while the trimmed model returns

$$P(\text{Failure of trimmed model}) = 1.80803714 \times 10^{-9}$$

The absolute error produced by trimming is 1.2×10^{-16} . The relative error (absolute error/true value) is 6.6×10^{-8} , which is about 1 part in 10 million. This example suggests that trimming significantly reduces model size and computational effort while producing an insignificant amount of error.

The trimming bound gives an upper bound on the absolute error produced by trimming. The trimming bound divided by the returned value for the trimmed model gives an upper bound for the relative error produced by trimming. In practice, if this upper bound for the relative error is small enough, then the trimmed model is acceptable. The amount of relative error that is acceptable varies with the application, but a relative error of 10 percent or less is usually considered acceptable.

This practice (of considering only the trimmed model and the trimming error bound) can be illustrated for this example. As will be shown subsequently, the upper bound for trimming error is

$$\text{TRMBND} = \theta\mu(e^{\theta T} - \theta T - 1)$$

where

θ maximum sum of the rates for the failure transitions leaving any state

μ largest average holding time for all recovery-mode states

T operating time

For this example,

$$\begin{aligned}\theta &= 3\lambda_P + 3\lambda_M + 3\lambda_B \\ &= (3 \times 10^{-4} + 15 \times 10^{-4} + 0.3 \times 10^{-4})/\text{hour} \\ &= 18.3 \times 10^{-4}/\text{hour} \\ \mu &= 1/F_M = 10^{-3} \text{ hour}\end{aligned}$$

since F_M is the slowest recovery rate, and $T = 1$ hour. The computed bound on the trimming error is

$$\text{TRMBND} = 3.066 \times 10^{-12}$$

This value of 3.066×10^{-12} is a bound on the absolute error produced by trimming. An upper bound for the relative error introduced by trimming is

$$\begin{aligned}\text{Relative error} &= \frac{\text{TRMBND}}{\text{Trimmed model result}} \\ &= (3.066 \times 10^{-12}) / (1.80803714 \times 10^{-9}) \\ &= 1.7 \times 10^{-3}\end{aligned}$$

which indicates that this model can be trimmed with negligible loss of accuracy. Note that the relative error for the upper bound of 1.7×10^{-3} is obtained without using the complete reliability model. The decision to use only the trimmed model has been made on the basis of the results from the trimmed model and the error bound for trimming.

The results for this example are typical for an application of trimming. The number of states in the reliability model is reduced by about half. The computational effort is reduced by about an order of magnitude. The actual error from trimming is insignificant. The derived error bound for trimming is much larger than the actual error, but the derived error bound is still small compared with the computed probability of failure for the system.

Description of Model Trimming and Statement of the Trimming Bound Theorem

A common approach to achieving reliability is to have three or four components perform a majority vote. When a component becomes faulty and disagrees with the majority, it is discarded from the system and replaced by a spare if a spare is available. There are two failure modes: (1) a coincident-fault failure when the voter is overwhelmed because a second component becomes faulty before a first faulty component can be removed and (2) an exhaustion-of-parts failure when the number of good components falls below a minimum level. Almost all fault-tolerant systems currently being considered are assemblages of subsystems each of which is a majority-voting system of the type described above.

For this class of systems, a reliability model has normal-operating states where all faulty components (if any) have been removed from the system, recovery-mode states where a faulty component has not yet been removed from the system, and absorbing states where the system has failed because of coincident faults or exhaustion of parts. We examine the recovery-mode states more closely. There are three types of transitions from a recovery-mode state: (1) system recovery, (2) failure of another component that causes immediate system failure (either coincident fault or exhaustion of parts), or (3) failure of another component that does not cause immediate system failure. Note that the third type of transition is not a transition to a system failure state. Model reduction by trimming is accomplished by removing all transitions of the third type (and their subsequent states) from the model.

Theorem: Suppose that

1. Components fail at a low constant rate.
2. Fault recovery depends only on the time since fault occurrence.
3. The system is an assemblage of subsystems, each subsystem achieving fault tolerance by a three-way or four-way majority vote.
4. All transitions to system failure are component failure transitions. (This assumption eliminates pathological cases.)

For each state i in the reliability model let

- θ_i sum of the component failure rates out of state i
- θ maximum value among the θ_i




Figure 3. General path in a semi-Markov reliability model.

μ_j average holding time in recovery-mode state R_j

μ maximum value among the μ_j

T system operating time

Then an error bound for model reduction by trimming is

$$\text{TRMBND} = \theta\mu(e^{\theta T} - \theta T - 1)$$

Derivation of the Trimming Bound

The error bound for model reduction by trimming is obtained from a theorem that places an upper bound on the probability of traversing a path in a semi-Markov reliability model by time T (refs. 4 and 5). In such a model, the component failures are assumed to occur at a low constant rate, and system recovery is allowed to be a fast, arbitrary distribution

that depends only on the time elapsed since component failure. A general path in such a semi-Markov reliability model is shown in figure 3. The global time independence of a semi-Markov process permits the rearrangement of states on the path for notational and computational convenience (refs. 4 and 5). In figure 3, small Greek letters represent slow constant-rate failure transitions, while capital roman letters represent fast system recovery transitions. The first line in figure 3 contains the states (A_k) with only slow constant-rate failure transitions (λ_k and γ_k). In the second line are states (B_i) where the successful (on-path) transitions are the fast recovery transitions ($F_{i,1}$) competing with slow constant-rate fault transitions (ϵ_i) and possibly other fast transitions (F_{i,b_i}). In the third line are the states (C_j) where the successful (on-path) transitions are the slow fault occurrences (α_j) competing against one or more recoveries (G_{j,c_j}) and possibly other fault transitions (β_j). For notation, let




Figure 4. Path diagram for the derivation of the trimming bound.

$P(F_i)$ probability that $F_{i,1}$ is successful
 $\mu(C_j)$ average holding time in state C_j
considering only recovery transitions

An upper bound for the probability of traversing the path in figure 3 by time T is

$$UB = \prod_{k=1}^{\ell} \frac{\lambda_k T}{k} \prod_{i=1}^m P(F_i) \prod_{j=1}^n \alpha_j \mu(C_j) \quad (1)$$

The general model for using the algebraic upper bound in equation (1) to derive the trimming bound is shown in figure 4. This general model displays all the paths to the possible system failure states in a reliability model for the class of system that we are considering. Since the model in figure 4 is potentially infinite, it includes transient faults and their potentially infinite occurrences.

The system starts in state A_0 which is a fault-free state. In this initial state, some component

failures can take the system immediately to system failure. These component failures are represented by the transition ϵ_0 to the system failure state X . Other components fail with rates $\alpha_1, \dots, \alpha_i$ and take the system to recovery-mode states R_1, \dots, R_i . From each of the R states, the diagram displays the three types of transitions out of a recovery-mode state discussed above. The ϵ transitions into X are the component failures that cause immediate system failure. The F transitions from recovery-mode states to fault-free states represent the possible system recovery actions. The γ transitions represent the component failures out of a recovery-mode state that do not cause immediate system failure.

After a γ transition, three simplifying assumptions are made about system behavior:

1. After a γ transition, the system is no longer able to remove failed components from the system.
2. After a γ transition, any other component failure causes immediate system failure.

3. This last transition (causing immediate system failure) occurs at rate θ which is the largest possible rate for a transition to a failure state, since θ is the maximum sum of the failure rates out of any state.

All these assumptions increase the computed probability of system failure (compared with the actual probability of system failure).

Returning to the main sequence of component failure and system recovery, the F recovery transitions out of the R states go to the fault-free B states where the cycle of component failure and system recovery begins again.

An upper bound is obtained for the probability of being in state Y_1 in figure 4 by considering all the paths to this failure state. One such path is the three-step transition from A_0 to R_i by α_i , from R_i to $V_{i,j}$ by $\gamma_{i,j}$, and from $V_{i,j}$ to Y_1 by θ . An upper bound for traversing this path by time T given by formula (1), the algebraic upper bound, is

$$UB = (\alpha_i T)(\gamma_{i,j} \mu) \left(\theta \frac{T}{2} \right) \quad (2)$$

where a slow failure transition competing with other failure transitions contributes the first factor, a slow failure transition competing with recovery transitions when the holding time in the recovery-mode state is less than or equal to μ contributes the second factor, and a second slow failure transition contributes the third factor. Summing over the fan of transitions from A_0 and the fan of transitions from the R_i 's gives

$$P(Y_1) \leq \sum_i \sum_j \left(\alpha_i \gamma_{i,j} \mu \theta \frac{T^2}{2!} \right) \quad (3)$$

$$P(Y_1) \leq \theta \mu \frac{T^2}{2!} \sum_i \alpha_i \left(\sum_j \gamma_{i,j} \right) \quad (4)$$

Since the sum of the failure rates out of any state is less than or equal to θ ,

$$P(Y_1) \leq \theta \mu \frac{T^2 \theta^2}{2!} \quad (5)$$

A typical path from A_0 to Y_2 has the five transitions from A_0 to R_i by α_i , from R_i to $B_{i,j}$ by $F_{i,j}$, from $B_{i,j}$ to $S_{i,j,k}$ by $\beta_{i,j,k}$, from $S_{i,j,k}$ to $V_{i,j,k,q}$ by $\gamma_{i,j,k,q}$, and from $V_{i,j,k,q}$ to Y_2 by θ . An upper bound

for traversing this path by time T is given by formula (1) as

$$UB = (\alpha_i T)[P(F_{i,j})] \left(\beta_{i,j,k} \frac{T}{2} \right) (\gamma_{i,j,k,q} \mu) \left(\theta \frac{T}{3} \right) \quad (6)$$

Summing over all the fans gives

$$P(Y_2) \leq \sum_i \sum_j \sum_k \sum_q \left[\alpha_i P(F_{i,j}) \beta_{i,j,k} \gamma_{i,j,k,q} \mu \theta \frac{T^3}{3!} \right] \quad (7)$$

$$P(Y_2) \leq \theta \mu \frac{T^3}{3!} \sum_i \alpha_i \left\{ \sum_j P(F_{i,j}) \times \left[\sum_k \beta_{i,j,k} \left(\sum_q \gamma_{i,j,k,q} \right) \right] \right\} \quad (8)$$

Since the sum of the failure rates is less than or equal to θ , and the sum of the probabilities for the recovery transitions $F_{i,j}$ is less than or equal to 1,

$$P(Y_2) \leq \theta \mu \frac{T^3 \theta^3}{3!} \quad (9)$$

In general,

$$P(Y_k) \leq \theta \mu \frac{T^{k+1} \theta^{k+1}}{(k+1)!} \quad (10)$$

Summing all these bounds for the Y_k 's gives a trimming bound of

$$\begin{aligned} \text{TRMBND} &\leq \sum_{k=1}^{\infty} P(Y_k) \\ &= \sum_{k=1}^{\infty} \theta \mu \frac{T^{k+1} \theta^{k+1}}{(k+1)!} \\ &= \theta \mu (e^{\theta T} - \theta T - 1) \end{aligned} \quad (11)$$

An Example With a Large Error Bound

A theorem on the error produced by trimming is necessary since trimming does not always have a negligible effect. Consider a system consisting of four reconfigurable fourplexes. Each fourplex removes itself from the system when the fourplex recovers from the second fault occurrence in that fourplex. The system fails by exhaustion of parts when all four fourplexes have removed themselves from the

system. A system coincident-fault failure occurs if any fourplex has a coincident-fault failure. In addition, all reconfiguration ceases if there are two faults present in two different fourplexes. In this case the system fails upon the occurrence of a third fault anywhere in the system. For a component failure rate of 10^{-4} /hour, a recovery rate of 10^3 /hour, and an operating time of 700 hours, the error bound for trimming is

$$\text{TRMBND} = 1.5 \times 10^{-6}$$

The computed probability of system failure using a trimmed model is

$$P(\text{Failure trimmed model}) = 7.05 \times 10^{-7}$$

The error bound for trimming is larger than the value returned by the trimmed model. Hence, the theory indicates that this model should not be trimmed. The computed probability of failure using the complete model is

$$P(\text{Failure complete model}) = 1.16 \times 10^{-6}$$

Trimming this model produces a significant error.

Concluding Remarks

This report has presented a method of model reduction called trimming and has derived an error bound for this method of reducing the number of states in a semi-Markov reliability model. The error bound uses only three parameters from the semi-Markov model: the maximum sum of rates for failure transitions leaving any state, the maximum average holding time for a recovery-mode state, and the operating time for the system. The error bound can be computed before any model generation takes place so that the modeler can decide immediately whether or not the model can be trimmed. The trimming has a precise and easy description which

makes it easy to include in a program that generates reliability models. This report has presented the simplest version of the error bound for trimming. Tighter bounds can be obtained by requesting more information about the system being modeled. For example, the current bound does not require any information about system recovery from multiple faults. Conducting the necessary experiments and including this information in the derivation of the error can produce a tighter bound. The price of the tighter bound is the cost of the experiments.

This method of model reduction is currently being developed as a feature of the automatic model generator called ASSIST.

NASA Langley Research Center
Hampton, VA 23665-5225
March 29, 1991

References

1. Johnson, Sally C.: *ASSIST User's Manual*. NASA TM-87735, 1986.
2. White, Allan L.; and Palumbo, Daniel L.: State Reduction for Semi-Markov Reliability Models. *Annual Reliability and Maintainability Symposium—1990 Proceedings*, IEEE Catalog No.: 90CH2804-3, Inst. of Electrical and Electronics Engineers, Inc., 1990, pp. 280-285.
3. Johnson, Sally C.: Reliability Analysis of Large, Complex Systems Using Assist. *A Collection of Technical Papers—AIAA/IEEE 8th Digital Avionics Systems Conference*, Oct. 1988, pp. 227-234. (Available as AIAA-88-3898-CP.)
4. Butler, Ricky W.; and White, Allan L.: *SURE Reliability Analysis—Program and Mathematics*. NASA TP-2764, 1988.
5. White, Allan L.: Reliability Estimation for Reconfigurable Systems With Fast Recovery. *Microelectron. & Reliab.*, vol. 26, no. 6, 1986, pp. 1111-1120.



National Aeronautics and
Space Administration

Report Documentation Page

1. Report No. NASA TP-3089	2. Government Accession No.	3. Recipient's Catalog No.	
4. Title and Subtitle Model Reduction by Trimming for a Class of Semi-Markov Reliability Models and the Corresponding Error Bound		5. Report Date May 1991	
		6. Performing Organization Code	
7. Author(s) Allan L. White and Daniel L. Palumbo		8. Performing Organization Report No. L-16862	
		10. Work Unit No. 505-66-21	
9. Performing Organization Name and Address NASA Langley Research Center Hampton, VA 23665-5225		11. Contract or Grant No.	
		13. Type of Report and Period Covered Technical Paper	
12. Sponsoring Agency Name and Address National Aeronautics and Space Administration Washington, DC 20546-0001		14. Sponsoring Agency Code	
15. Supplementary Notes An early version of this material was presented at the 1990 Annual Reliability and Maintainability Symposium and appears in the proceedings.			
16. Abstract Semi-Markov processes have proved to be an effective and convenient tool for constructing models of systems that achieve reliability by redundancy and reconfiguration. These models are able to depict complex system architectures and to capture the dynamics of fault arrival and system recovery. A disadvantage of this approach is that the models can be extremely large, which poses both a model construction and a computational problem. Techniques are needed to reduce the model size. Because these systems are used in critical applications where failure can be expensive, there must be an analytically derived bound for the error produced by the model reduction technique. This report presents a model reduction technique called trimming that can be applied to a popular class of systems. Automatic model generation programs have been written to help the reliability analyst produce models of complex systems. This method (trimming) is easy to implement and its error bound easy to compute. Hence, the method lends itself to inclusion in an automatic model generator.			
17. Key Words (Suggested by Author(s)) Reliability estimation Reconfigurable systems Semi-Markov models Error bounds Model reduction Automatic model generation		18. Distribution Statement Unclassified—Unlimited Subject Category 65	
19. Security Classif. (of this report) Unclassified	20. Security Classif. (of this page) Unclassified	21. No. of Pages 9	22. Price A02

